arXiv:1703.07841v1 [cs.NE] 22 Mar 2017

Classification-based RNN machine translation using GRUs

Ri Wang
rtwang@uwaterloo.ca

Maysum Panju
mhpanju@uwaterloo.ca

Mahmood Gohari
mgohari@uwaterloo.ca

Abstract

We report the results of our classification-based machine translation model, built upon the framework of a recurrent neural network using gated recurrent units. Unlike other RNN models that attempt to maximize the overall conditional log probability of sentences against sentences, our model focuses a classification approach of estimating the conditional probability of the next word given the input sequence. This simpler approach using GRUs was hoped to be comparable with more complicated RNN models, but achievements in this implementation were modest and there remains a lot of room for improving this classification approach.

1 Introduction

Machine translation is an increasingly important field of study, as the need for natural language translation is rapidly rising in a world that is growing increasingly multilingual and interconnected. At the same time, machine translation is an incredibly difficult task whose problems span a range of linguistic fields and is difficult even for human translators. Previous efforts have been centered on statistical machine translation, but the continuous and exponential increase in computing power and advances in neural network training have enabled researchers to use recurrent neural networks to approach this problem.

Currently, neural networks such as LSTMs and bidirectional RNNs are trained for this purpose. However, methods described in many of these research papers are relatively complex and frequently non-intuitive. The approach taken by feed-forward neural network researchers, however, involves using the neural network as a classifier to predict probabilities, and is a simpler and more attractive framework to build upon. In this work, we attempt to remove the common constraint in feed-forward networks that require input sequences to have a fixed-length. Instead, we propose to build a relatively simple GRU neural network model, which both allows variable-length input and is able to act as a classifier applied to the task of machine translation.

2 Background and Related Work

The approach for translation using recurrent neural networks was mainly based on the model given by Sutskever et. al. [12] of recursively inputting output into input to generate a translation. In their paper, they trained pairs of long short-term memory models (LSTM's): one LSTM served as a feature extractor to convert a variable-length sentence into a fixed-length vector such that sentences with similar meanings are mapped close together, and the other LSTM takes the word vectors as hidden states and proceeds to output words one at a time, incrementally building up the translation. It was not entirely clear how this pair of LSTM training was done in the context of back-propagation, and we decided to focus on their alternative model of having a single RNN to perform both feature extraction and translation. The model we used was also very closely related to that of Cho et al. [1], which introduced the GRU and applied it as an RNN encoder-decoder machine translation model as an alternative to LSTM.

However, the classification approach taken here is closer related to previous work done by Schwenk [10], with feed-forward neural networks that take in a fixed input sentence and output the probability of a sequence of words occurring next, as well as the work of Devlin et al., who only predicted the next word [3]. The difference in our approach is that

the GRU used here can take in a variable-length sentence, and is much more flexible than the models by Schwenk and Devlin et al., which were restricted to only accepting input of a pre-specified fixed length.

3 Basic Framework of the Learning Models

The key ingredient for training a machine translation model is a large corpus of sentences in the source language that has been correctly translated into sentences in the target language. For a powerful translator, this training corpus must be very large and the translations must be perfectly accurate.

There are other specifications and requirements that need to be met as well. The model requires both source and target languages to draw from fixed vocabulary sets; in this work, we restrict both the input and output vocabulary sets to each consist of 80,002 distinct tokens, representing the 80,000 most frequently appearing words in the training set for each language, along with 2 additional tokens added to represent the not-found token and the end-of-sentence token. To decrease the number of unique words, all words were converted into lowercase, but were neither stemmed nor lemmatized. This is because grammatical inflections in words are very important to overall translation quality, as they carry heavy semantic value that needs to be preserved during translation.

Prior to use in the model, words were converted into word vectors using a standard continuous word-to-vector embedding [7]. The translation model would eventually be trained to recognize that semantically similar words should be weighted to carry similar representations, and using word embeddings that already indicate semantic similarity from the outset should reduce a lot of training effort on the part of the translation model. Furthermore, we expect higher quality word embedding from specialized representation models, such as GloVe [8], than any word embedding built into the translation neural network such as the one used by Cho et al. [1], because they are trained on a much larger dataset and their linguistic capabilities have been proven to be very strong.

The training was based on stochastic gradient descent with mini-batches of 128 sentence pairs each. To take full advantage of GPU parallelization, batches were prepared so that within each batch B_i, all of the source-language input sequences consisted of the same number of tokens F_i, and all of the target-language output sequences consisted of the same number of tokens E_i, although it was not necessary that F_i be the same as E_i for any given batch. This condition was enforced to prevent part of the batch dropping out of the back-propagation gradient calculations midway, in the event that the some sequences reached completion before other sequences in the batch. If there were not enough sentence pairs in the training set corresponding to a given pair of lengths (F_i, E_i) to fill a minibatch of size 128, then the entire batch B_i was dropped from the dataset altogether; this measure filtered out a large number of small batches, which had been observed to learn much less than full-size batches, despite taking almost the same amount of time to train. Additionally, both F_i and E_i were hard-capped with a maximum of 50, because the neural network would have trouble handling translations between lengthy sentences of longer than 50 tokens in either the source or target language.

4 Model

In this section, we describe the procedure that the deep learning model is used to perform machine translation, once the preprocessing and setup from Section 3 has been implemented.

4.1 Automatic Machine Translation

After the input text in the source language has been represented as sequences of word embedding vectors of fixed length, our translation model works as follows:

For a given sequence of input word vectors $(x_1, \ldots, x_F)$ encoding a sentence in the source language, as well as the sequence of output word vectors $(y_1, \ldots, y_k)$ encoding the tokens in the translated sentence in the target language so far, the neural network estimates the conditional probability of the word vector representation of the next word in the translated sequence y_{k+1}. Thus, part of the input sequence to the neural network would be a complete sequence of source-language word vectors and the remainder of the input would be a sequence in progress consisting of target-language word vectors.

The output vector of this neural network is a vector of length 80,002, with components corresponding to each of the words in the target language vocabulary, along with the two special tokens representing not-found and end-of-sentence.

The entry at position i in this vector, for $1 \leq i \leq 80,002$, is a real-valued number between 0 and 1 representing the probability that the ith word in the target language vocabulary set, v_i, is the next word in the translation sequence.

If the neural network is treated as a function f, then its operation is defined as:

$$f((x_1, \ldots, x_F), (y_1, \ldots, y_k)) = [P(y_{k+1} = v_i | (x_1, \ldots, x_F), (y_1, \ldots, y_k))]_{i=1,\ldots,80,002}$$

As the translation is built up word-by-word, the neural network can be treated as a classifier that takes in a sequence of input vectors and predicts the "class", or vocabulary word, that should appear next in the sequence.

Finally, to calculate the conditional probability of obtaining the output translation sentence $y_1, \ldots, y_E$ given the input sentence word sequence $x_1, \ldots, x_F$ for the purpose of translation, the probability may be decomposed into a product of conditional probabilities that the neural network is able to compute:

$$P((y_1, \ldots, y_E) | (x_1, \ldots, x_F) = \prod_{k=0}^{E-1} P(y_{k+1} | (x_1, \ldots, x_F), (y_1, \ldots, y_k))$$

The predicted translation is then the sentence with the highest overall conditional probability, where the probabilities may be obtained by calculating each term in the above product using the neural network. At each stage, the input is the source sentence sequence $(x_0, \ldots, x_F)$ and the incrementally built-up target sequence $(y_0, \ldots, y_k)$.

4.2 The neural network model: GRUs

The main contribution of deep learning to the task of automatic machine translation, in the approach taken by this work, is the computation of the conditional probabilities for the next translated word given the progress of the translation so far. This approach has been applied using various specializations of neural networks for the task, including long short-term memory networks [5] and bidirectional recurrent neural networks [9]. In this work, we explore the possibility of using gated recurrent units in the model.

A gated recurrent unit (GRU) was initially proposed by Cho et. al. [1] for their translation task. The idea behind this approach is to keep the current content at each step in the neural network, and add the new contents to that, while forgetting any past information that no longer adds additional information to the present state. Unlike the LSTM approach that define a separate memory cell, the GRU resets a gate and stores only the relevant new content to the unit. The two approaches of GRU and LSTM in handling variable length inputs are presented for contrast in Figure 1.

Given an input sequence $x = (x_1, x_2, ..., x_T)$, the standard RNN computes the activation by

$$h_t = a(Wx_t + Uh_{t-1})$$

where a is a smooth activation function such as hyperbolic tangent function. The GRU computes an update gate as

$$z_t^j = \sigma^j(W_z x_t + U_z h_{t-1})$$

This update gate is the sum of current state at t and the recent state which controls the amount of previous memory to be kept. In the LSTM, the update gate has an additional component that is the memory cell at time t. Similar to the update gate, a reset gate is computed based on the current input vector and hidden state, but with different weights:

$$r_t^j = \sigma^j(W_v x_t + U_v h_{t-1})$$

Now the activation is computed as

$$\tilde{h_t}^j = \tanh^j(Wx_t + U(r_t \odot h_{t-1}))$$

where $\odot$ represents pointwise multiplication. If the reset gate is zero, then the previous state will be ignored and only the current input is used for updating. The final output at time t is computed as a weighted average of $\tilde{h}_t$ and h_t

$$h_t = z_t \odot h_{t-1} + (1 - z_t) \odot \tilde{h}_t$$

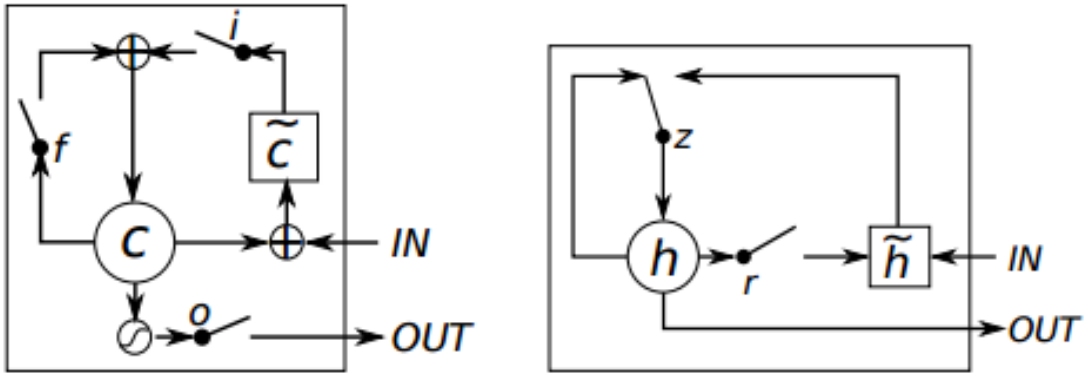

(a) Long Short-Term Memory (b) Gated Recurrent Unit

Figure 1: Long-Short Term Memory (left) and Gated Recurrent Units (right) gating. c and $\tilde{c}$ stand for the memory cell and the new memory cell content in LSTM. r and z are the reset and update gates, and h and $\tilde{h}$ are the activation and the candidate activation in GRU. Figure copied from Chung, J, et al.: Empirical evaluation of gated recurrent neural networks on sequence modeling [2].

5 Training Methods and Implementation

For this work, we made use of the publicly available Europarl[1] corpus [6], with French as the source language and English as the target language. This dataset contains about two million instances of English and French sentence pairs across a large variety of sentence lengths. Translation was done from French to English because this report is intended for an English-speaking audience, and a model with English-language output allows us to assess the quality of the translation directly and independently. Furthermore, this approach distinguishes our work from that of many other papers in the field, where English is frequently taken to be the source language.

For the English word vector embedding, we used the pre-trained model made publicly available by the researchers for GloVe on their website at http://nlp.stanford.edu/projects/glove/. For French, however, a pre-trained embedding was not conveniently available, so a model was built using GloVe as part of this work, trained using the entire corpus of French-language Wikipedia. Although this did provide a workable word embedding for use in the translation model, the size of the training set for these French word embeddings was limited to only the half a billion tokens that were available, in contrast with the 42 billion tokens that were used to train the provided English word embeddings. In both cases, the representations that were returned belong to a 300-dimensional vector space.

Training was performed using batched stochastic gradient descent, where each batch contained about 128 sentences. The model was set to have 4 layers, in light of the success observed by Sutskever et. al [12]; however, the number of nodes in the present model is set at 500 per layer, in contrast with the 1000 used by the earlier work, in order to speed up training and reduce potential overfitting.

In this work, the focus was on training the neural network to work as an effective classifier, which is the key component of the translation framework laid out in Section 4.1. There were two ways to approach this in the context of automatic translation, as outlined in the following sections.

5.1 Method 1

In this training method, the focus is entirely upon building a strong classifier under the belief that everything has gone perfectly up to the current position in the translation. Given a sequence of French input word vectors $(x_0, \ldots, x_F)$ and the corresponding sequence of *correct* English translation word vectors $(y_0, \ldots, y_E)$, this method generates predicted

[1]The European Parliament dataset, restricted largely to political topics and certain styles of speaking, is not representative of speech as a whole, which means the model might not be able to generalize well under other circumstances. However, it is not easy to obtain high quality translation data, and we have to make do with what is available.

translation words $\hat{y_k}$ one at a time, but uses the ground-truth translations instead of its own predictions every time it looks for the next word in the sequence:

1. input: $((x_1, \ldots, x_F), ())$; output: $\hat{y_1}$; expect: y_1
2. input: $((x_1, \ldots, x_F), (y_1))$; output: $\hat{y_2}$; expect: y_2
   ...
3. input: $((x_1, \ldots, x_F), (y_1, \ldots, y_{E-1}))$; output: $\hat{y_E}$; expect: y_E

This model is built upon high quality input and should therefore act a powerful classifier. However, since it is only trained upon flawless input data, it might encounter difficulties in practice if any word happens to be mistranslated part-way through a sequence, as the model does not have experience correcting itself from a partially incorrect translation sequence.

5.2 Method 2

The focus of the second training method is to build a classifier that is able to handle robustly correct itself from intermediate translation errors. As in the previous approach, this method generates predicted translation words $\hat{y_k}$ one at a time, but now uses the model's own predicted translation, rather than the ground-truth, to predict the next translated word in the output sequence:

1. input: $((x_1, \ldots, x_F), ())$; output: $\hat{y_1}$; expect: y_1
2. input: $((x_1, \ldots, x_F), (\hat{y_1}))$; output: $\hat{y_2}$; expect: y_2
   ...
3. input: $((x_1, \ldots, x_F), (\hat{y_1}, \ldots, \hat{y_{E-1}}))$; output: $\hat{y_E}$; expect: y_E

This model is much more practical and robust than the previous one, as it is trained on its own output, much like in the situations where the model is applied in practice; as a result, there should be a much smaller difference between test error and training error, and the model will perform better in real applications.

The major disadvantage with this approach is that training is a much more difficult and time-consuming task. For much of the early part of training, the output predicted words will be close to random and the neural network will have to learn how to use this random noise to predict the next correct word. Not only would each step in the prediction process be slower during the training period, but the required number of training iterations would rise by orders of magnitude, compared with the optimistic approach of Method 1.

5.3 Technical Details

Some of the specific decisions made regarding the implementation of model training for this work are described below.

1. Intial weights were set using the Glorot initialization [4] with weights sampled from the uniform distribution. This is designed to set weights that are neither too small nor too large, so as to avoid vanishing or exploding signals.
2. We did not expect many issues with vanishing gradients, since the GRU-based model is generally able to safely handle this problem. To avoid the unlikely risk of exploding gradients, a hard clip for gradient values was kept at 100.
3. The entire collection of batches is shuffled at the start of every training epoch. Training then proceeds stochastically through the sequence of batches. There were about 6000 batches in total and a single epoch took 5-6 hours using a single Nvidia 980 Ti processor.
4. The training error for each word prediction in the sequence was computed by taking the cross entropy between the predicted vector of conditional probabilities vector produced by the neural network, and the ground-truth vector which has a value of 1 at the index of the correct next word and 0 for all other indices.
5. Nesterov momentum [11] was used during the gradient descent process in an attempt to speed up the training procedure.

6 Results

In general, the trained neural network for this work did not perform as well as other state-of-the-art models, but was able to achieve a modest success considering the constrained time and resources that were available.[2] Of the two methods described in Section 5, the second method is believed to provide a stronger and more robust translation model. However, in practice, the training procedure was prohibitively slow as the model could not identify useful patterns in translation quickly from the original randomly-generated translations. As a result, only the first method was successfully implemented and tested.

The results of the translation model can be qualitatively evaluated by looking at some of the automatically translated sentences, and comparing them with the correct English translations. A sample of some translations can be seen in Table 1. For comparison, the translation offered by Google Translate is also provided.

Through the present model, translation quality quickly deteriorates for longer sentences, even those within the training set. The results on some of the shorter sentences, on the other hand, are reasonable and sometimes even creative. However, it is clear that this model is not powerful enough to fully translate by itself, achieving a validation test score of only 22.9% – that is, given an input sequence of words in French and a partial sequence of words in the English translation so far, the probability of identifying the next word in the translation correctly is a little more than 1 in 5.

A possible reason for the struggling performance of this model overall is that the neural network was trained to estimate the conditional probability given a correct sequence of inputs, but the translation process builds upon predicted output words that may not be correct. Once the predicted word sequence goes "off track", it no longer has a good context to translate from, and having never learned how to go back to the correct sequence, the model will understandably face difficulty. It should be noted that if the model is asked to predict the next word when given an input sequence of words in French and a partial sequence of words in the *correct* English translation, the probability of identifying the next word in the translation correctly is nearly doubled, at about 40.3%.[3]

7 Recommendations and Conclusions

The GRU-based model presented in this work showed early signs of potential for leading to a powerful automatic translator. Although the resulting translations were not always of the highest quality, there is almost always a visible link to the intended translation, and there is frequently some coherency within the overall prediction. As a starting attempt at a new approach with tight constraints on time and resources, the present success may even be considered quite promising.

In future work, it would be interesting to try using a larger training set, more nodes per hidden layer in the neural network, and a larger dataset for training the French language word vectors. These adjustments would certainly amplify the time required for training, but it is possible that they may lead to a more capable translator. Based on the solitary model examined as part of this work, it not clear what the effects that many of the parameters, such as number of nodes and the size of the vocabulary, would have on the ability of the resulting translator. This would certainly present an interesting avenue of future research.

8 Acknowledgments

The authors would like to extend grateful acknowledgments to Dr. Ghodsi for teaching the course and creating an environment of enthusiasm for learning, both for the students and our algorithms. Thank you as well for taking the time to mark our final projects during the winter holiday.

[2] It is unfortunate that despite repeated and earnest consultation with Sharcnet support services, technical issues prevented the research team from obtaining the full benefit of Sharcnet resources in time to complete this work.

[3] Perhaps 40% is not particularly good either, but machine translation is a difficult problem, and it is comforting to note that even Google Translate has issues with translating some sentences as well.

English Translated Sentence	English Correct Sentence	Google Translate	Set
am much more space . 1	i ve got plenty of room	I have a lot of space.	Test
few things are being used for the journey	a light raincoat is ideal for the trip	a light raincoat is ideal for travel.	Test
am delighted to see our next meeting	i look forward to our next meeting	I look forward to our next meeting.	Test
president on a point of order	madam president on a point of order	Madam Chair c is a procedural motion.	Training
you send a doctor to a doctor	will you send someone go get a doctor	are you send someone to fetch a doctor.	Test
please please call for this silence	please rise then for this minute silence	I invite you to stand for this minute of silence.	Training
he was born it was very difficult to speak . NF	since he was dressed in black he looked like a priest	as he was dressed in black he looked like a priest.	Test
will have a painful job but he did not have his surprise	she told him a joke but he didnt think it was funny	she told him a joke but he did not find it funny.	Test
have had a debate on this subject in the next part session in the debate . NF . NF .	you have requested a debate on this subject in the course of the next few days during this part session	you wanted a debate on this in the coming days during this session.	Training
declare resumed the session of the european parliament which was held on thursday 19 october which i have had to say and i am sure you will confirm that my meeting was good news .	i declare resumed the session of the european parliament adjourned on friday 17 december 1999 and i would like once again to wish you a happy new year in the hope that you enjoyed a pleasant festive period	resumed session I declare resumed the session of the European Parliament adjourned on Friday 17 December and I renew my vows in the hope that you had a good holiday.	Training

Table 1: Translated phrases obtained using our model, in contrast with the correct translation and the result obtained from Google Translate. "NF" implies that the model predicted a word outside of the vocabulary, or a potential sentence break.

References

[1] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.

[2] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.

[3] Jacob Devlin, Rabih Zbib, Zhongqiang Huang, Thomas Lamar, Richard Schwartz, and John Makhoul. Fast and robust neural network joint models for statistical machine translation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1370–1380, 2014.

[4] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *International conference on artificial intelligence and statistics*, pages 249–256, 2010.

[5] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[6] Philipp Koehn. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86. Citeseer, 2005.

[7] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.

[8] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. *Proceedings of the Empiricial Methods in Natural Language Processing (EMNLP 2014)*, 12:1532–1543, 2014.

[9] Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. *Signal Processing, IEEE Transactions on*, 45(11):2673–2681, 1997.

[10] Holger Schwenk. Continuous space translation models for phrase-based statistical machine translation. In *COLING (Posters)*, pages 1071–1080, 2012.

[11] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th international conference on machine learning (ICML-13)*, pages 1139–1147, 2013.

[12] Ilya Sutskever, Oriol Vinyals, and Quoc VV Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.